

1. Summary of Thursday's Session

A. We briefly reviewed the contents of the Annotator's Guide that was distributed earlier, the overview of SALAMI, and oriented ourselves in Sonic Visualizer. Everything should be covered in the following documents:

http://www.music.mcgill.ca/~jordan/salami/annotator_guide.pdf

<http://salami.music.mcgill.ca/>

<http://sonicvisualiser.org/doc/reference/1.7/en/index.html>

B. We reviewed two relatively simple examples: the Beatles' "Think For Yourself" and Stan Getz's "Dear Old Stockholm." Then we had you try out the system for yourself by annotating a couple harder pieces: "Nocturne No. 3 Op. 51, from Belshazzar's Feast" performed by Sarah Chang, and "Solo Song with Lyre Accompaniment" by Doogus Idris. Following each annotation exercise, we compared results and had a discussion about what works best in different situations. Audio files and annotations for all these pieces are downloadable here:

[not available]

From the discussion, it sounded like extra examples would be a great help to people, so there are a few more examples in that batch: one each from the broad popular, classical, and world collections.

To view any of the examples, open the audio file in Sonic Visualizer. You can do this through the menu (File > Import Audio File) or by dragging and dropping the audio file into a blank session. Overlaying the annotation file works the same way: either use the menu (File > Import Annotation Layer), or drag and drop the text file into the session. You'll then have to choose the following options:

1. Each row specifies **a point in time**
2. The first column contains **time, in seconds**

2. Homework

We've given you 6 pieces to annotate as homework, and we ask for you to complete your annotations for these pieces by next Friday, May 14th. You can download the files from here:

[not available]

At around 30 minutes per track, we're hoping you can complete them in roughly 3 hours. But please keep track of how much time you spend working on each song, and let us know—we want to see how long it actually takes. Remember that you will be paid for this time.

When you've finished annotating each song, save the Sonic Visualizer session (for your own records) and export the annotation layer to a text file. Each song has a single word

as its filename, so save that text file as “yourname-songname”. For example, if I were to export my annotations for the Haydn Symphony designated “miracle,” I’d name the file “smith-miracle.txt”.

When you’ve finished annotating all the songs, please email them to me. (You can zip the into a single folder if you like.) Remember to submit them by May 14th.

3. Annotations

Recall that there are three annotation layers: letters, functions, and instruments. The format rules for each layer are specified below. All of the provided examples conform to this format.

Letters: every segment should be annotated with either a letter (A, B, C, etc.) or a marker (X, =) which indicates a subsection. Different letters should indicate different musical ideas or phrases.

Functions: every segment that belongs to a section with a particular function (e.g., all the letters that are in the “verse” section) should be marked with that label.

- In the Stan Getz example (Dear Old Stockholm), note that since the opening AABA section is all part of the “head,” so each part of it has been labelled with “head.”
- In the Doogus Idris example (Solo Song), note that there are sections which aren’t marked with any function label. This is appropriate if the function is ambiguous, unknown.

Instruments: sections which contain a leading instrument should be marked as such, using opening and closing tags. For instance, the boundary marking the onset of a singer’s voice should have the label “(vocal”, and the last section containing voice should be marked “voice)”. If the voice is only present on a single section, you may write “(voice)”.

Bonus notation notes:

- Most pieces start and end with a period of silence. Both of these should be marked with the label “silence.”
- The very end of the piece should be marked as a boundary and labelled “end.”
- In any boundary label, you can include comments after a number sign (#). Note that *everything* after the # will be ignored, so don’t write:
A #might be a verse... (piano
- It is perfectly acceptable for zero or two instruments to be leading. As for more than two instruments leading... without making a strict proclamation, we can say that that case seems unlikely -- certainly it stretches the definition of “leading.” But of course there’s lots of unlikely music!

4. Tips

Using Sonic Visualizer:

Here is a work flow to annotate songs in Sonic Visualizer that works for me. The relevant hot-keys are mentioned in brackets.

1. Load the song into Sonic Visualizer.
2. Immediately press the semicolon key (;) to mark the beginning of the track.
3. Zoom in (up arrow) and use the pencil tool (4) to manually mark a boundary at the first onset of sound.
4. Listen to the whole song (spacebar) and make marks (;) at any and all times that you think sound like boundaries. Don't worry about missing pick-ups, since boundaries should be placed on downbeats.
5. If there's a period of silence at the end of the song, mark where it starts. When the song ends, press semicolon again to mark the final boundary.
6. When listening to the song, you can now quickly navigate it by skipping to the previous/next boundary (page up/down).
7. Bring up the list of boundaries you've made (e).
8. Listen through the song again, labelling sections as you go. (If you're a fast typer, you can do this during playback. Each row will be highlighted as it is played. Press tab until the third column is highlighted, and then start typing.)
9. While doing this, you might want to pause periodically (space bar) to adjust the boundaries using the edit tool (3).

After you've made a first pass on the labels, each song will work differently. Sometimes the structure is obvious, and it's just a matter of putting in function labels and then sorting out the leading instruments. Sometimes, if the original boundaries were hard to estimate, you'll want to jump back and forth between two versions of the same section to make sure that your labelling is consistent. In less straightforward situations, you'll find yourself re-listening to and re-annotating certain portions.

When you're done, it's helpful to give the whole piece a final look: start listening from the beginning again, skipping through the sections (page down) and reading the labels you've given. Check that the formatting is correct and that nothing's been missed.

Nice hot-keys:

1-5	selects the different tools (move, select, edit, etc.)
↑	zoom in
↓	zoom out
g	adds a spectrogram
;	add a "time instant"—your best friend!
e	brings up time instants Data Editor (so that you can add labels)
apple-z	undo—your other best friend!
page up	skip to previous time instant
page down	skip to next time instant
escape	unselect

Advice on Letters:

As we discussed on Thursday, the definition of what each letter is meant to represent is fuzzy: they should each designate a different musical idea, or phrase, or subject. But the characterization of that musical idea will depend on the context: in a rock song, where the same chords are repeated throughout the whole song, slight differences in melody or rhythm could deserve different labels. But in a classical piece with a much more through-composed design, assigning each melodic idea with a new letter would get you to Z pretty quickly, so you'll need to approach the labelling more abstractly. (The strictness with which you use X, =, or the prime symbol will be similarly adjustable.)

For instance, in the "Think For Yourself" example, each letter designates a very strict harmonic and melodic chunk. But in the "Solo Song" (Doogus Idris) example, the letter C designates a looser melodic idea, although it still stands apart as a separate musical idea from either the A or B sections. In the "Le Grand Tango" example, since the texture is very improvisational, some very different-sounding sections which use the same harmonic or rhythmic idea are designated with the same letter (for instance, the A and H sections). But in the context of the whole piece, which sprawls relatively widely, it is important to provide an annotation that is informative and succinct.

Remember that you're aiming to limit yourself to the letters A-E. It should be rare that you need to use more than 6 letters, and perhaps just as rare that you will need only one letter.

Advice on Functions:

When the project starts in earnest, we will try to standardize our vocabulary for function labels somewhat. For now you're free to use whatever labels you feel best suit the music at hand, but while doing so, keep in mind the end goal of a standardized vocabulary.

As proposed on Thursday, hierarchical labelling is permitted if it makes you feel more comfortable. (You may even find it helps you organize your thoughts better!) To do this, simply add a second time instants layer (under the Layer menu). You can select all (apple-A) and then copy (apple-C) paste (apple-V) the time instants into the new layer, and then delete the ones you don't need. Hierarchical labelling was used in one of the provided examples ("Le Grand Tango"), an exceptional case which was deemed to truly consist of three consecutive pieces. (It's thus also an example of using letters well beyond E.)

Good luck!